Congreso de Software IBM 2010

Construyendo un planeta más inteligente

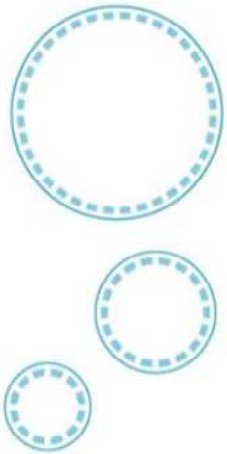

IBMAGINA

Un planeta más inteligente necesita software más inteligente

Carlos Guardia Rivas

Executive IT Specialist - IBM SWG
SPGI IT/Specialist Profession Leader
Member of the SPGI Technical Expert Council (TEC)

IBMAGINA
Un planeta más inteligente necesita software más inteligente

Gestión de la Información:

Datawarehouse y System Analytics en System Z

IBMAGINA

Un planeta más inteligente necesita software más inteligente

Data volume is exploding…

Volume of Digital Data

Variety of Information

Velocity of Decision Making

As of December 2008, the global monthly Internet traffic is estimated to be **5 to 8 exabytes**.

As of May 2009, the size of the World's total Digital content has been roughly estimated to be 500 billion gigabytes, or **500 exabytes.**

By 2013, annual global IP traffic will reach **two-thirds of a zettabyte or 667 exabytes.**

Internet video will generate over **18 exabytes** per month in 2013

Source: Wikipedia 3/2010

Yet businesses still struggle

Lack of Insight

One in three managers frequently make critical decisions without the information they need

Inefficient Access

One in two don't have access to the information across their organization needed to do their jobs

Inability to Predict

Three in four business leaders say more predictive information would drive better decisions

Source: IBM Institute for Business Value

IBMAGINA

Un planeta más inteligente necesita software más inteligente

The foremost issue facing enterprises today is the waste, inaccuracy and volume of missed opportunities that stem from information raging out of control

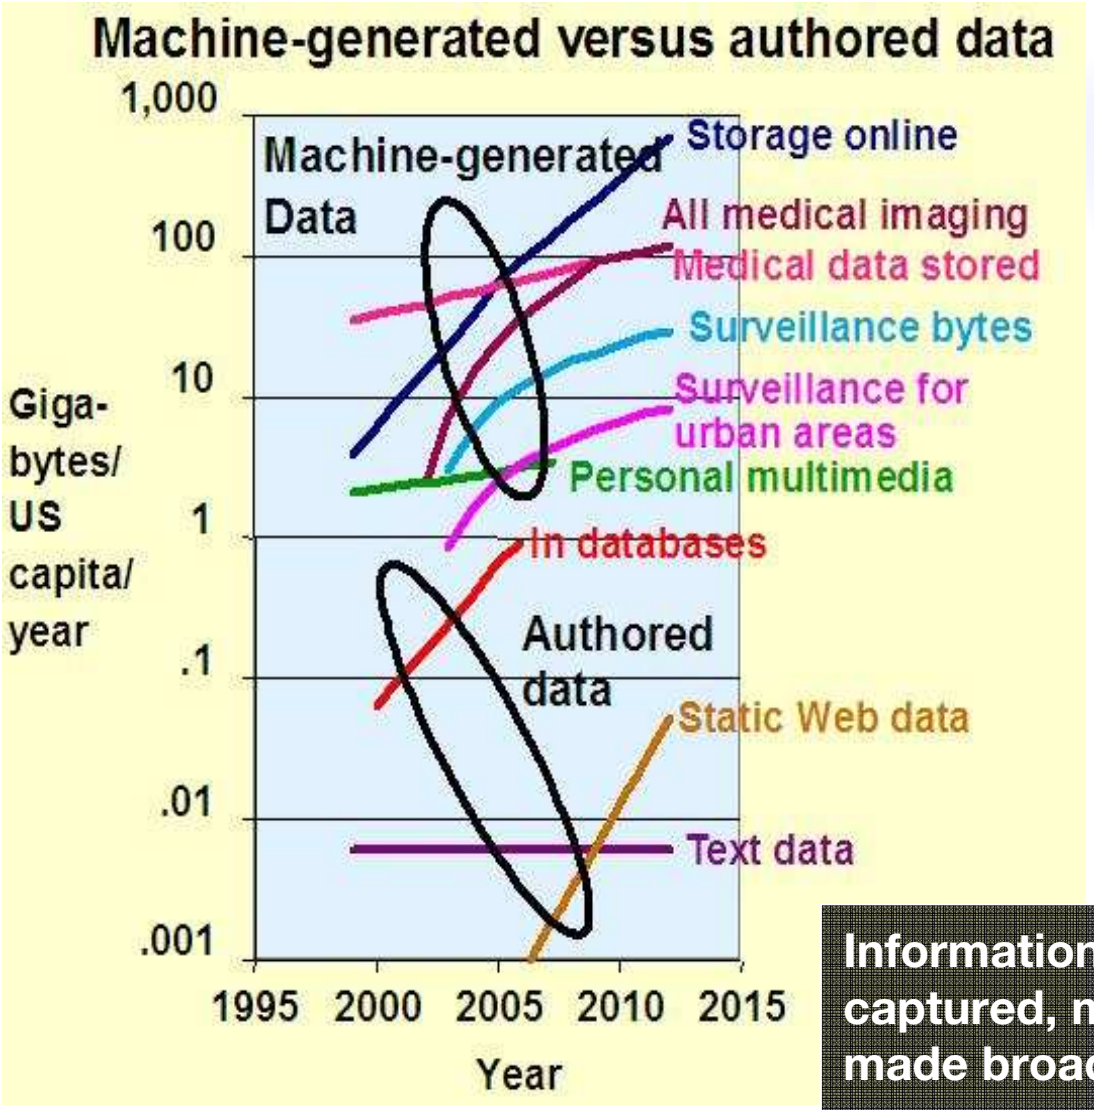

Volume of Digital Data

- 57% CAGR for enterprise data through 2010
- Machine generated data : Sensors, RFID, GPS..

Variety of Information

- 80% of new data growth is unstructured
- Emails, images, audio, video..

Velocity of Decision Making

- 42% of managers say they use wrong information at least once per week
- Rapidly changing business climate

Information is not being effectively captured, managed, analyzed and made broadly available to everyone who needs it

IBM IMAGINA
Un cliente más inteligente necesita software más inteligente

Business Intelligence Becoming Mission-Critical

Past	Today
Point in Time Business Intelligence	*Right-time Business Intelligence*
Batch Data Warehousing	*Active Data Warehousing*
Warehouse and Transaction Systems are uncoupled	*Continuous feedback into business processes*
Self-contained Historical Data Warehouses	*Information Integration with other data sources*
Latency in development and deployment of BI apps	*Time to Value is CRITICAL*

IBMAGINA
Un planeta más inteligente necesita software más inteligente

DataWarehousing trends

- **High requirement for system, platform and data** security
- **Desire to optimize & leverage** existing z infrastructure and skills
- **Minimize complex and costly** data movement
- **Need help managing high growth of '3 Vs" –** Volume, Variety, Velocity
- **Increasing demands for sophisticated analysis with real time operational data — BI is becoming** *mission critical*

Dynamic Warehousing
A New Approach to Leveraging Information

Information On Demand to Optimize Real-Time Processes

Dynamic Warehousing

OLAP & Data Mining to Understand Why and Recommend Future Action

Query & Reporting to Understand What Happened

Traditional Data Warehousing

MAGINA

Planeta más inteligente necesita software más inteligente

Data Warehousing on DB2 for z/OS – What is driving this?

- **Many System z customers already use DB2 for z/OS for warehouse and BI**
  - IBM is responding to customer demand with new DB2 features, new software offerings and improved hardware performance and efficiency

- **The System z platform offers superior Total Cost of Ownership**
  - Customers want to leverage their System z infrastructure to do more with what they are already using successfully
    - TCO can be reduced through the utilization of existing processors, people, practices
    - TCO may also be achieved through a consolidation approach

- **New BI trends map well to the strengths of DB2 for z/OS and System z**
  - The distinction is blurring between warehouse and OLTP databases based on new trends such as Dynamic Warehouse and Operational BI, driving:
    - The need for increased reliability, availability, security, and compliance in a DWH DBMS
    - The need for very current warehouse data and/or collocation of warehouse and operational data

- **Specialty processors and the new z10 provide additional ways to optimize TCO**
  - zIIPs and IFLs are driving down hardware and software costs; DWH/BI can make excellent use of these processors, ultimately driving TCO advantages
  - The new processors are delivering excellent speeds

IBMAGINA
Un planeta más inteligente necesita software más inteligente

DB2 for z/OS features that support DWHing

DB2 V8

- **64-Bit Addressability**
- 2000 byte index keys
- **MQT's**
- Multi-Row Operations
- 225 way table joins
- In-Memory Workfiles
- Automatic Space Allocation
- **Non-Uniform Distribution Statistics on Non-Indexed Columns**
- **Parallel Sorting**
- **Data Partitioned Secondary Indexes**
- 2MB SQL Statements

DB2 9

- **Partition by growth**
- **Index Compression**
- **Index on expression**
- **Dynamic index ANDing**
- RRF new row internal structure for VARCHAR
- Fast delete of all the rows in a partition
- Skipping uncommitted inserted/updated qualifying rows
- **Histogram Statistics**
- SQL enhancements
  - INTERSECT, EXCEPT, cultural sort, caseless comparisons, RANK, DENSE_RANK, …
- **DB2 VUE**

DB2 10

New!

- Numerous optimizer enhancements
  - Minimize materialization and size of intermediate results
  - **Hash access path**
  - Parallel index update at insert
  - Faster single row retrievals
- **Index include columns**
- **Auto statistics collection**
- **Compress 'on the fly'**
- Enhanced query parallelism for improved performance
  - Remove query parallelism restrictions
  - In-memory techniques for faster query performance
  - Advanced query acceleration techniques
- **System z synergy**:
  - **Improved cpu cache performance**
  - **Exploit new h/w instructions**
  - **Utilize z10 1MB page size**
- **Moving sum, moving average**
- **Advanced query acceleration**

DB2 for z/OS for DWHing

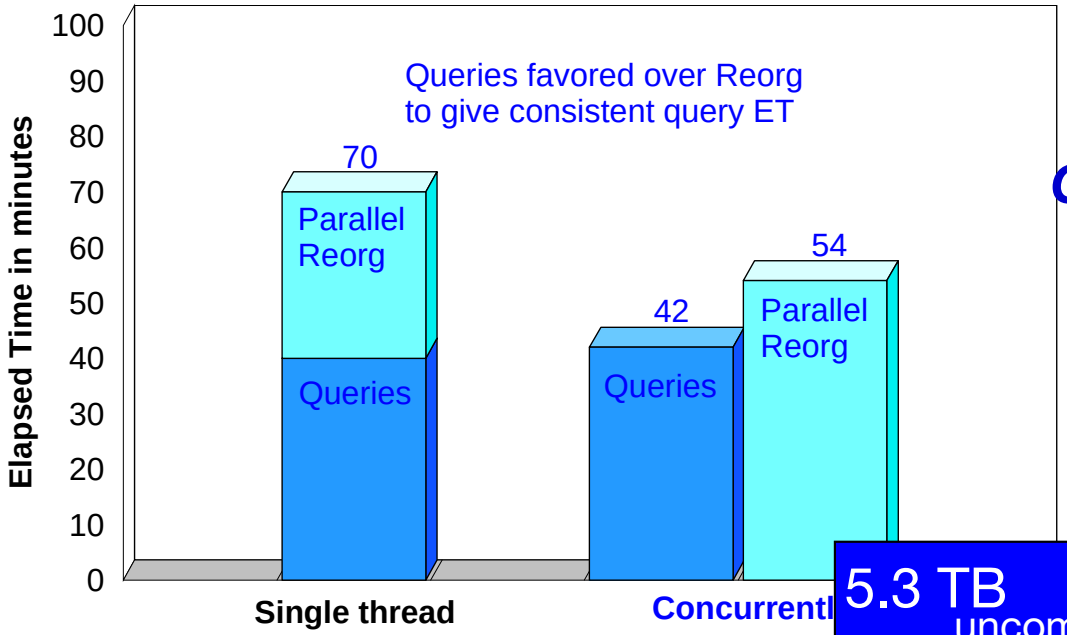

Concurrent workload management

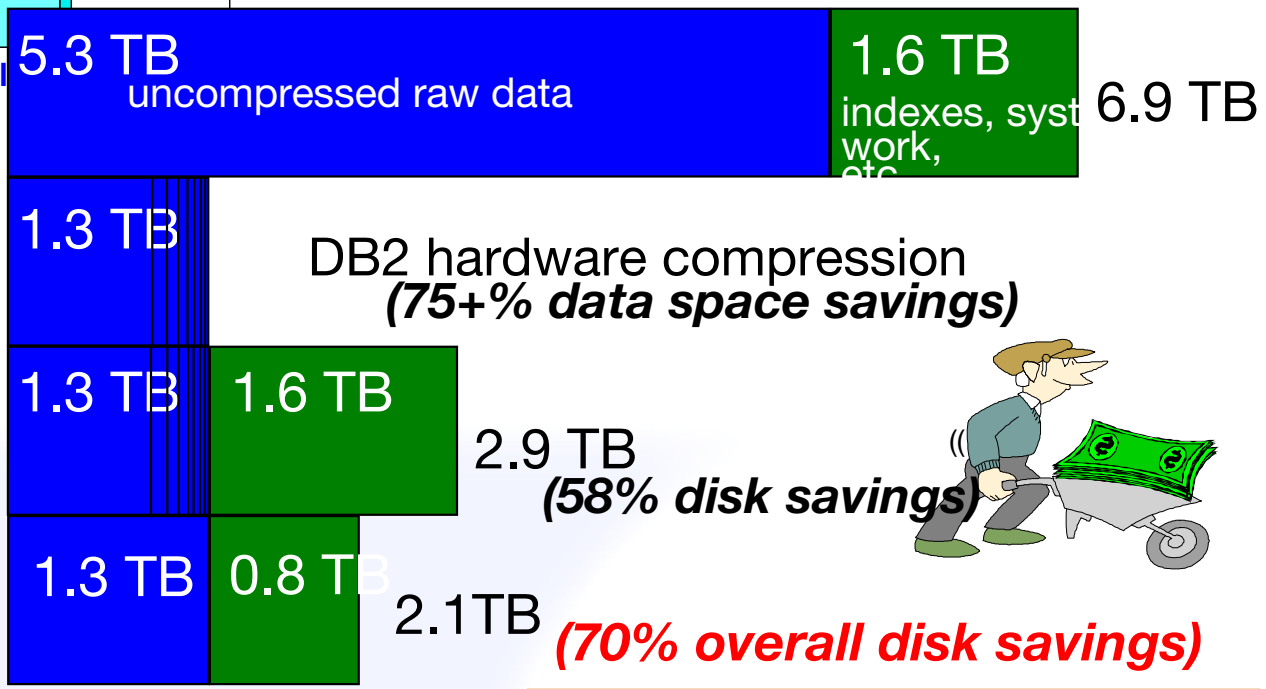

High compression ratios

DB2 for z/OS workload management

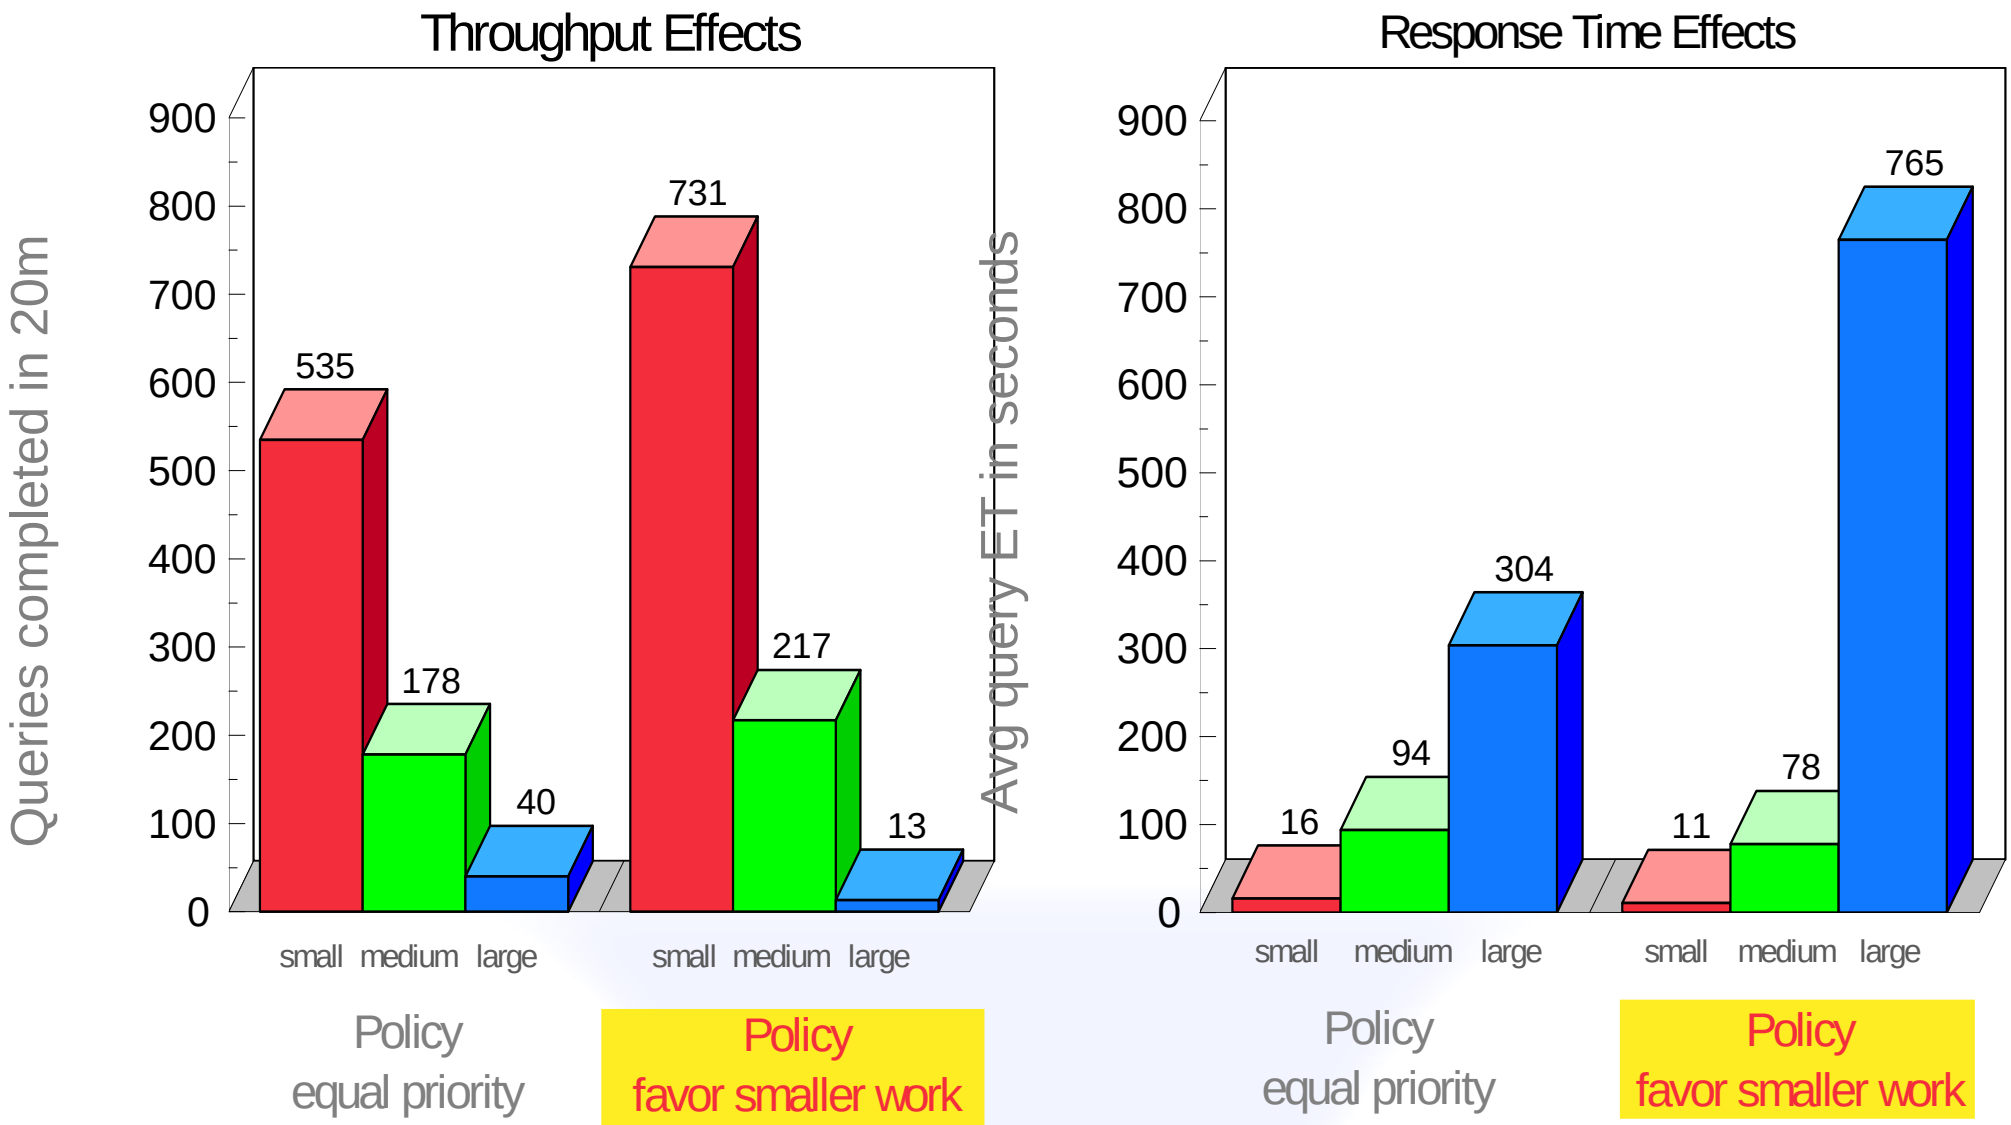

Information-Led Transformation for System z

More new capabilities delivered in the past 3 years than at any point in the mainframe's history

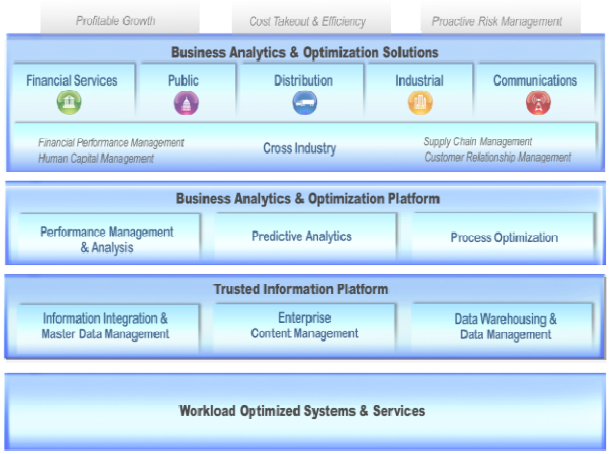

Smart Analytics Optimizer for z/OS
DB2 end–to-end monitoring
DB2 and IMS Active-Active for z/OS
DB2 10 for z/OS
DB2 Sort for z/OS
Master Data Management for z/OS
Cognos 8 BI for z/OS
2010
Entity Analytics for System z
IMS Tools Solution Packs for z/OS
IMS 11 for z/OS
Analytics
Content Integrator for z/OS
InfoSphere Warehouse for System z
Information-Led Transformation
FileNet Content Manager for System z
Cognos 8 BI for System z
2009
InfoSphere MDM Server for System z
Business Intelligence
Optim Data Governance for System z
Information Server for System z
Information Agenda
Content Manager 8.4
2008
Information On Demand
Data Warehousing

DB2 software
InfoSphere software
FileNet software
Cognos software

The IBM Data Warehouse Solution on System z

Ultimate Consolidation Opportunity

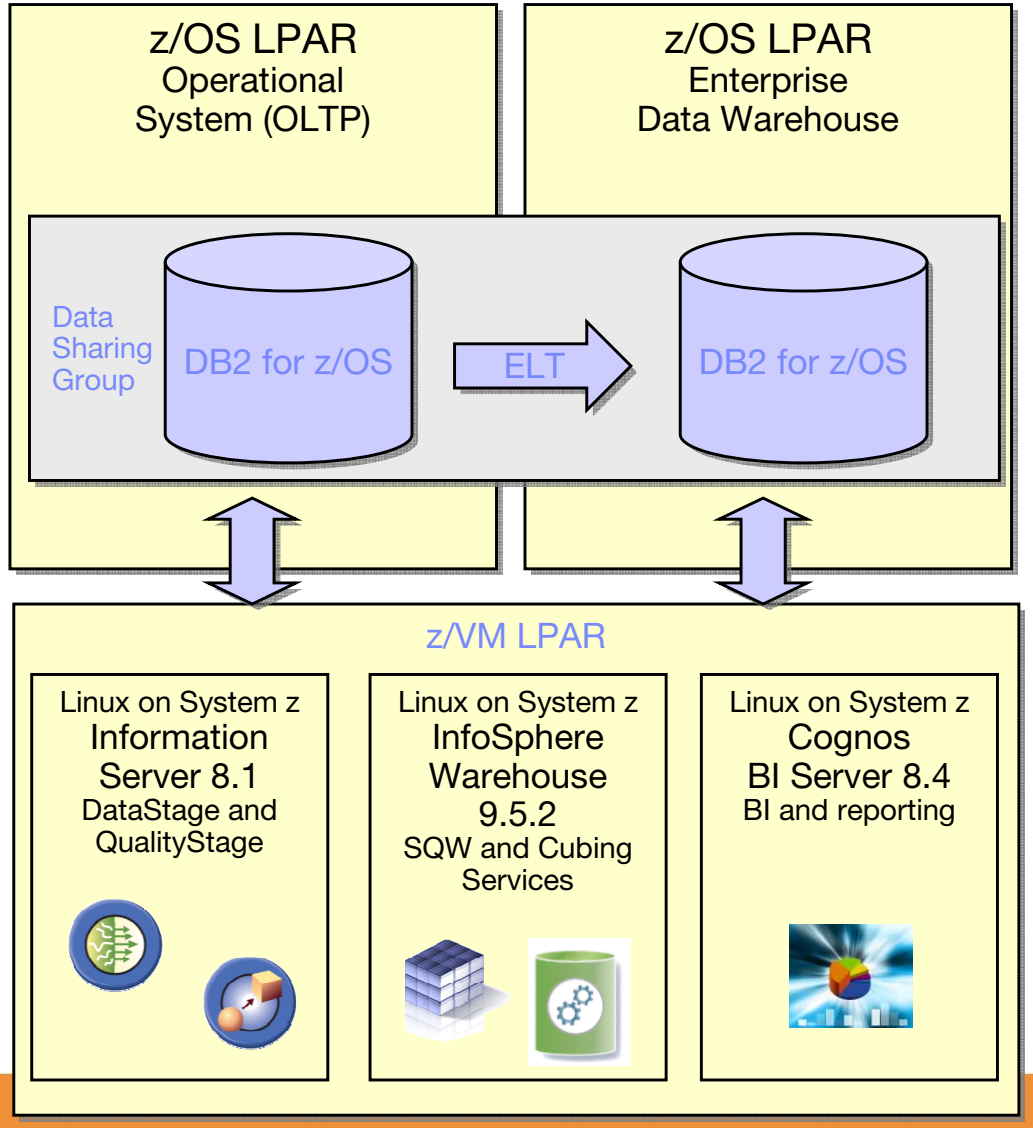

- Consolidation of mission-critical data on System z
- Ability to leverage existing environment, high availability, backup and governance procedures as well as skills
- Efficient data movement within a data sharing group (no network)
- Performance and TCO improvements through cubing services (data marts) and DB2 enhancements
- Complex transformations and data quality driven from Linux on System z with Information Server

IBMAGINA
Un planeta más inteligente necesita software más inteligente

IBM Smart Analytics Optimizer

Adding Industry Leading Performance

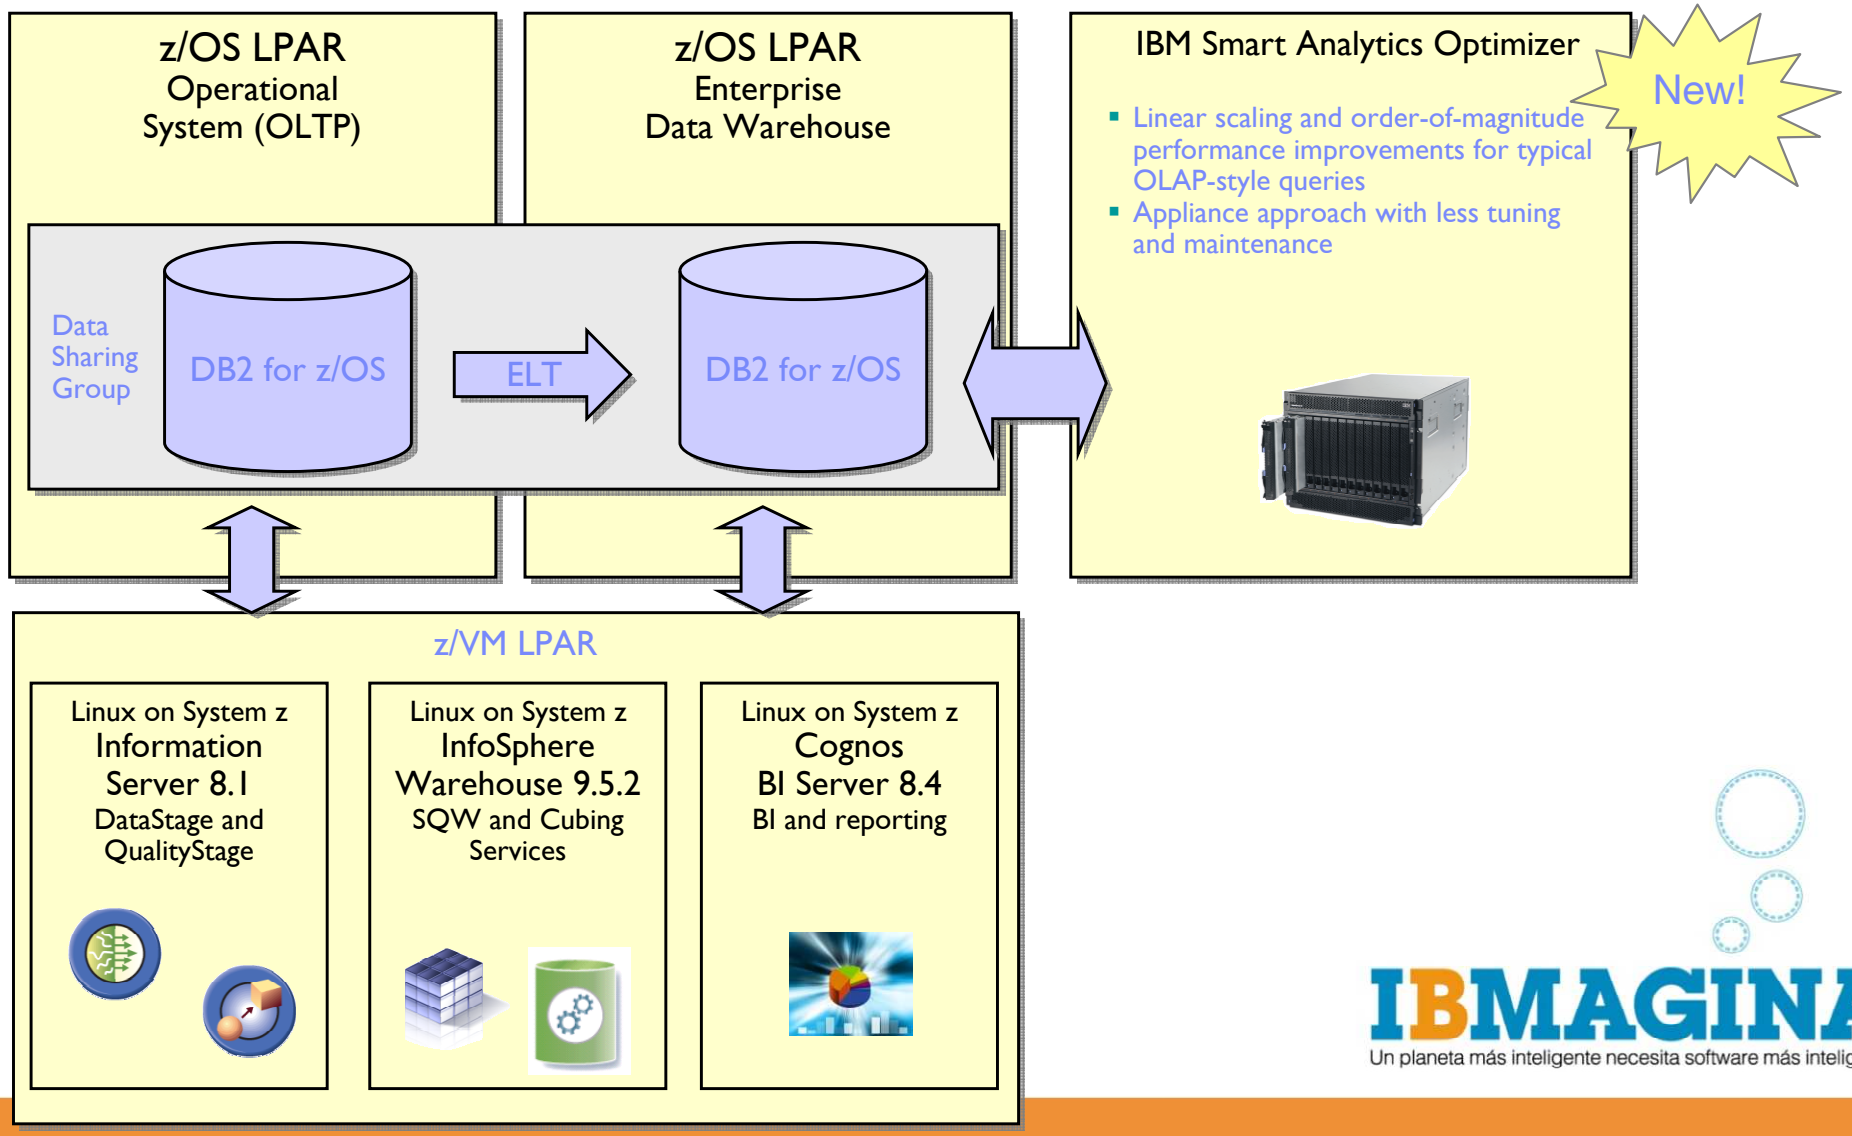

IBMAGINA

Un planeta más inteligente necesita software más inteligente

The IBM Smart Analytics Optimizer Solution

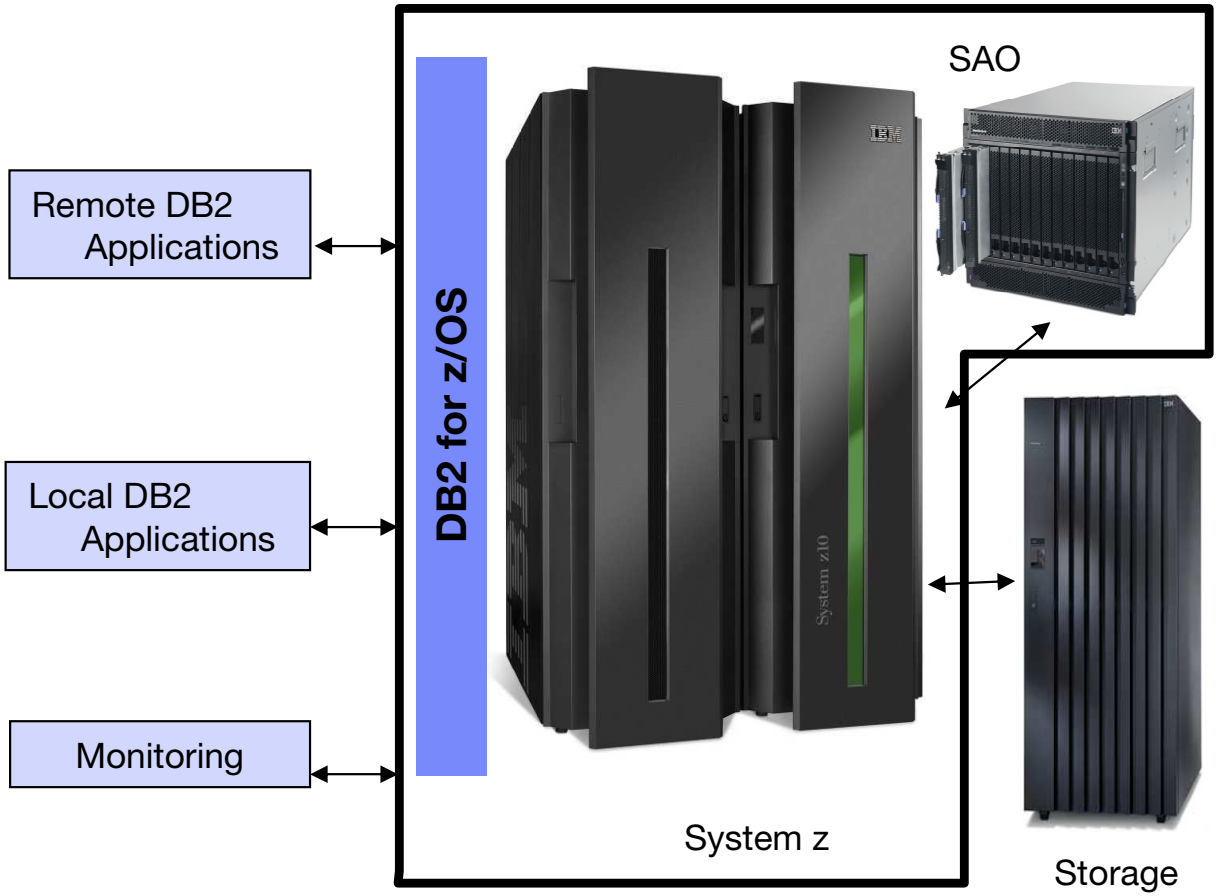

- A special purpose, network attached x86 blades system
  - Offload typical DW queries from traditional database server to the accelerator
  - Based on research prototype called BLINK
- No changes to the applications
  - Applications continue to attach to DB2.
  - DB2 transparently to the applications exploits the accelerator when applicable query needs to be executed
- Improving performance of typical DW queries 5-10 times
- Achieving linear scaling with the number of CPUs
- Reducing need for tedious tuning of DB2 (MQTs, indexes, etc.)
- Significantly improved price/performance and TCO as a combined effect of:
  - Offloading very CPU intensive operations from System z
  - Using commodity hardware
  - Reduced DBA effort for tuning offloaded queries

Query Response Times SAO faster in all cases w/ predictable response

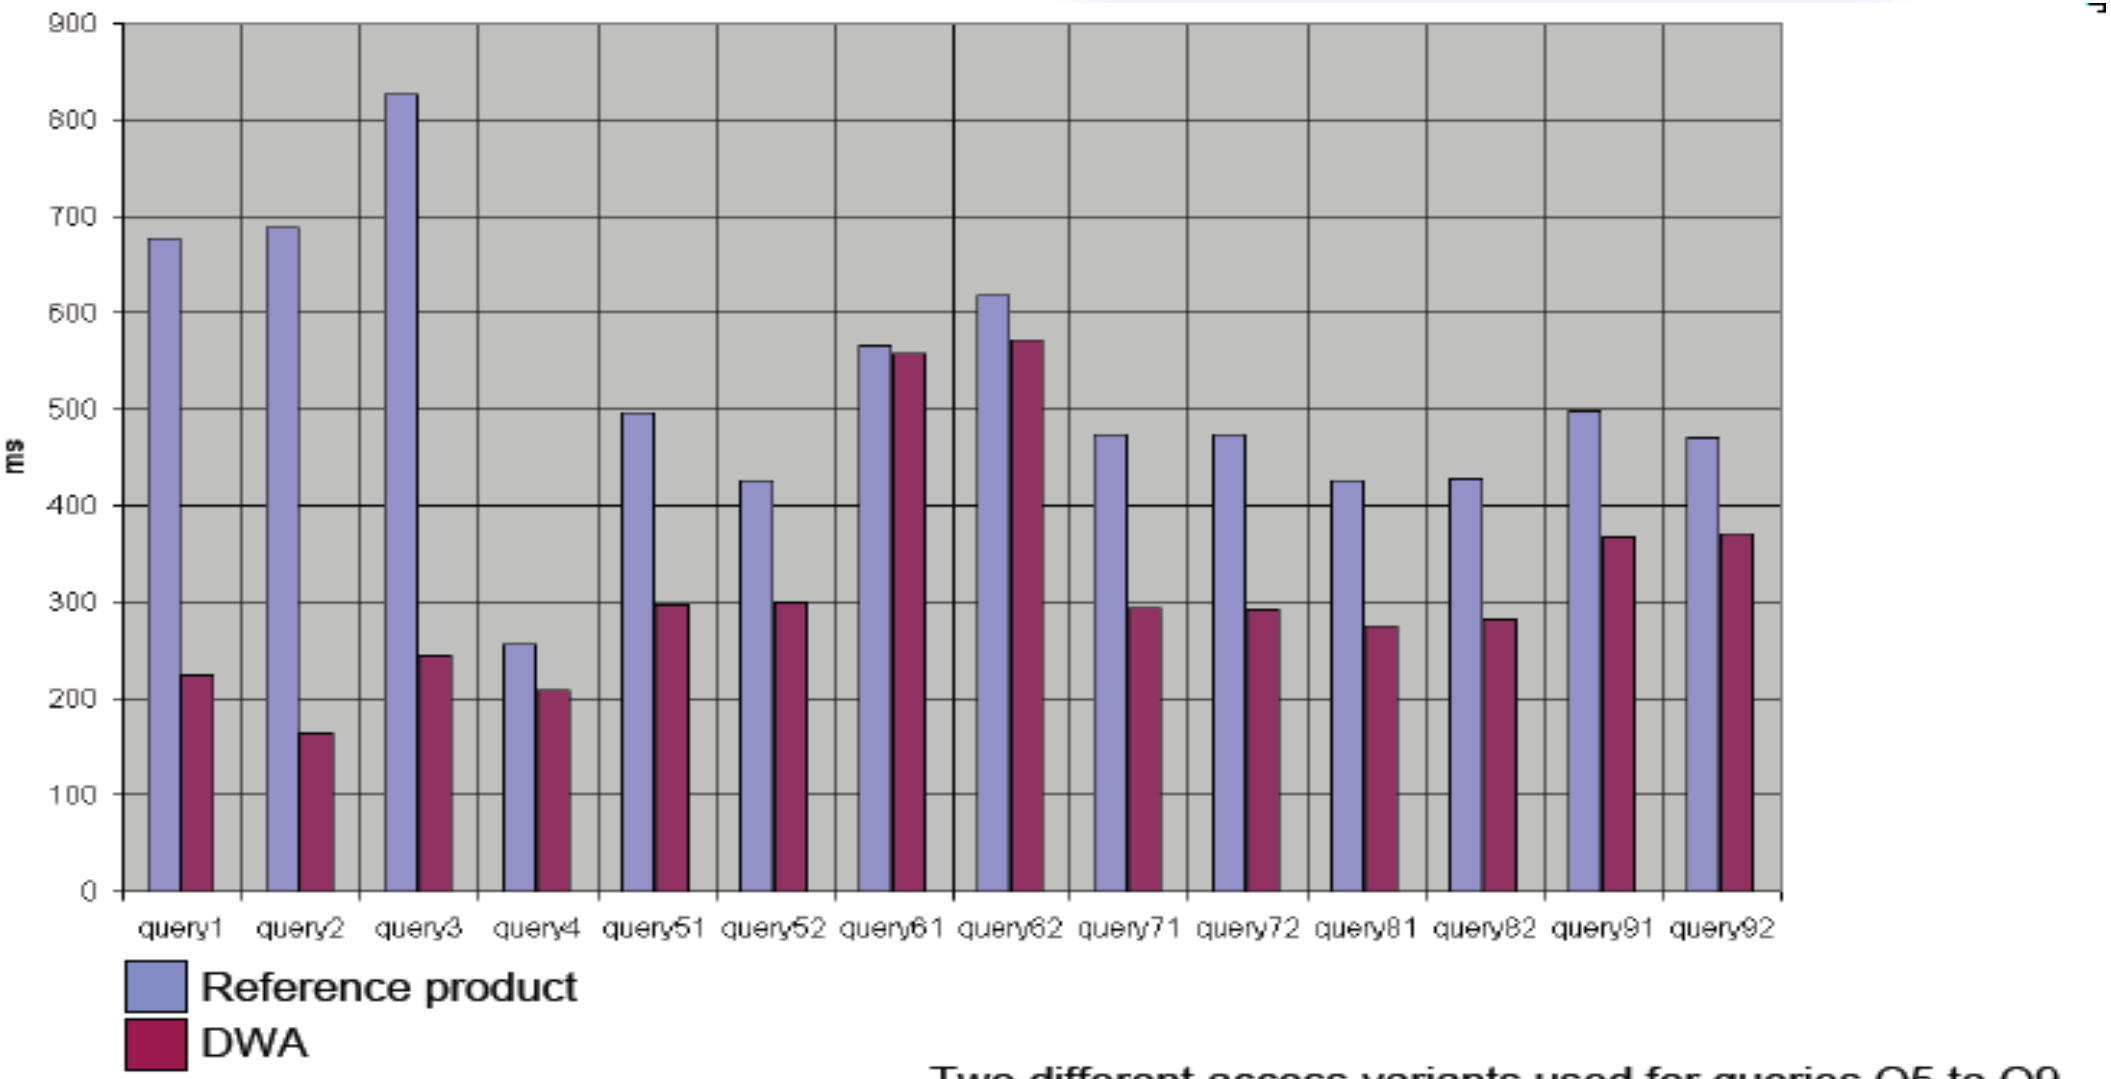

Two different access variants used for queries Q5 to Q9.

IBM Smart Analytic System 9600

- An integrated hardware, software and services offering that enables enterprise customers to quickly and cost effectively deploy game changing analytics across their business.
- What does it include?
  - ✓DB2 for z/OS Value Unit Edition (Primary, with option for MLC)
  - ✓DB2 Utilities Suite (Default in, Ability to opt-out)
  - ✓InfoSphere Warehouse for Linux on System z
  - ✓Cognos 8 BI for Linux on System z (Default in, Ability to opt-out)
  - ✓z/OS Operating System Stack
  - ✓Implementation Services
  - ✓Dedicated z10 server or bolt-on LPARs to existing server
    - ✓Includes general processor, zIIP and IFL engines to meet capacity requirements

Optional Software

- Tivoli OMEGAMON for DB2 Performance Expert
- DB2 Connect
- Tivoli Directory Server
- InfoSphere Information Server
  - All components eligible
- InfoSphere Replication Server
  - Q-Rep, CDC and Event Publisher eligible
- InfoSphere Federation Server plus Classic Federation on System z
- SPSS
- Tivoli ITCAM, ITUAM
- Cognos Now! For Linux on System z
- Cognos Blueprints for Healthcare, Banking and others…
- BI User on-boarding application (as proposed for Smart Analytics Cloud)

IBM Smart Analytic System 9600: High Value Dynamic Warehousing

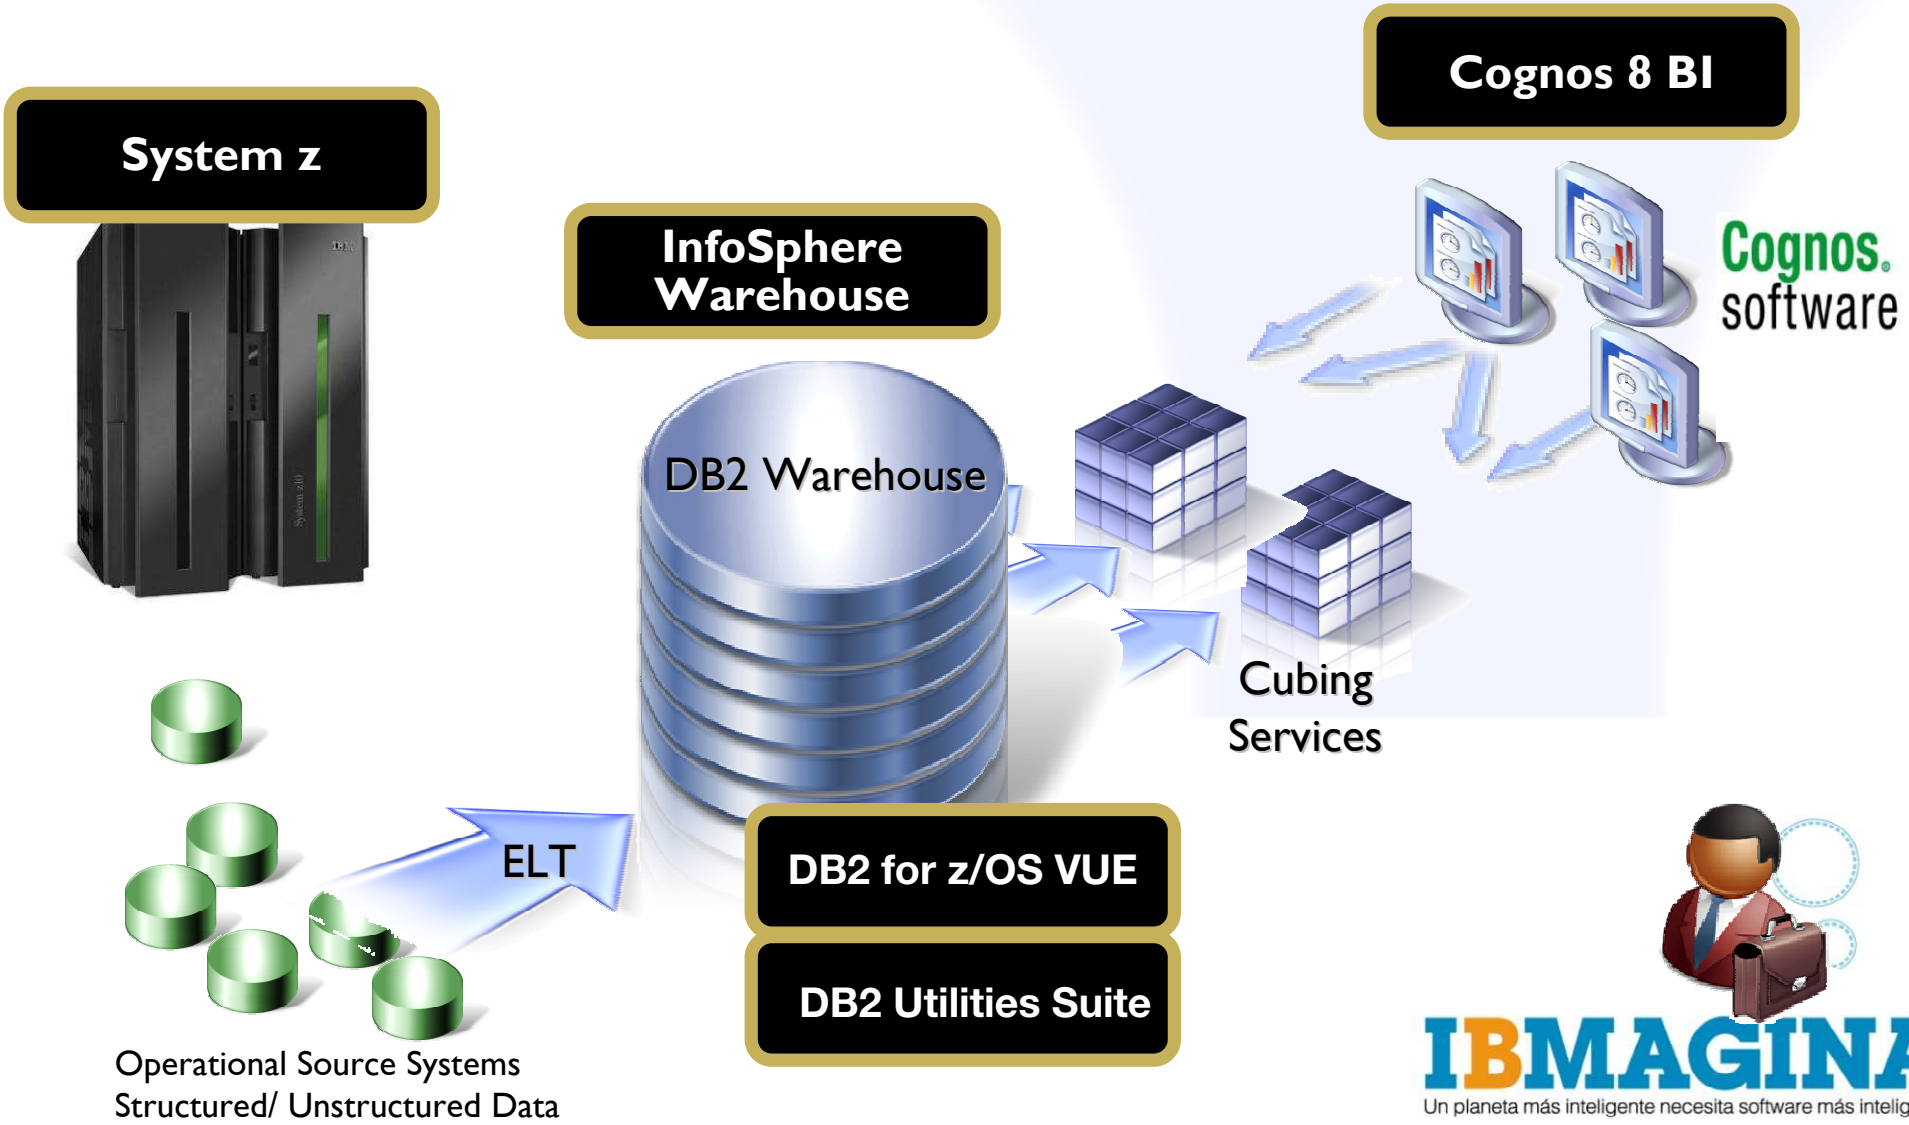

Cubling services: OLAP

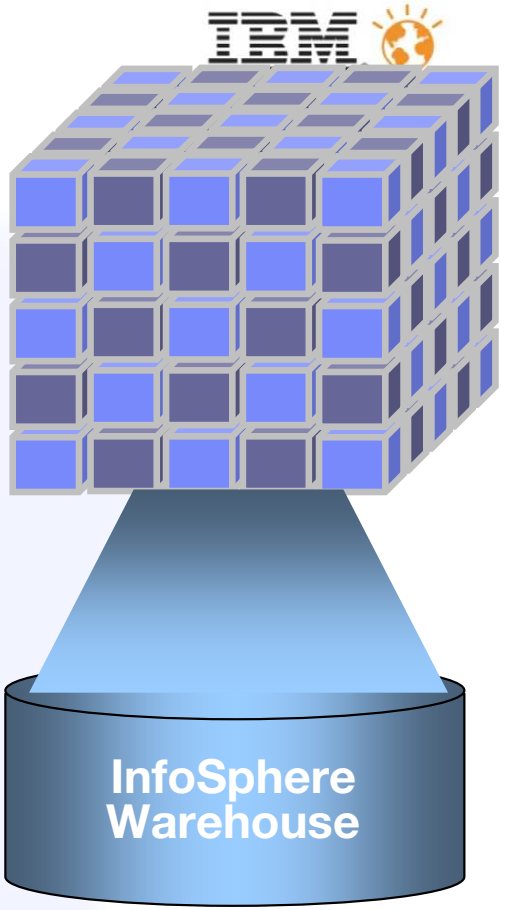

Cubing Services is a Multidimensional Analysis Server that enables OLAP applications access to Terabyte data volumes via industry standard OLAP connectivity

Benefits

- Uses industry standards OLAP Interfaces for access by a wide variety of tools for presentation and reporting
- Empowers users with ad hoc access to business information.
  - What is the profitability for Product A across the Branches X,Y,Z?
- Speed of thought access to OLAP data managed in InfoSphere Warehouse. Thus for your InfoSphere Warehouse OLAP data there is:
  - OLAP and SQL shared access to the same information
  - Single point of Management
  - Single point of Maintenance
  - Single point of Performance Tuning
- Enables Access to up to 1 TB of base OLAP Data

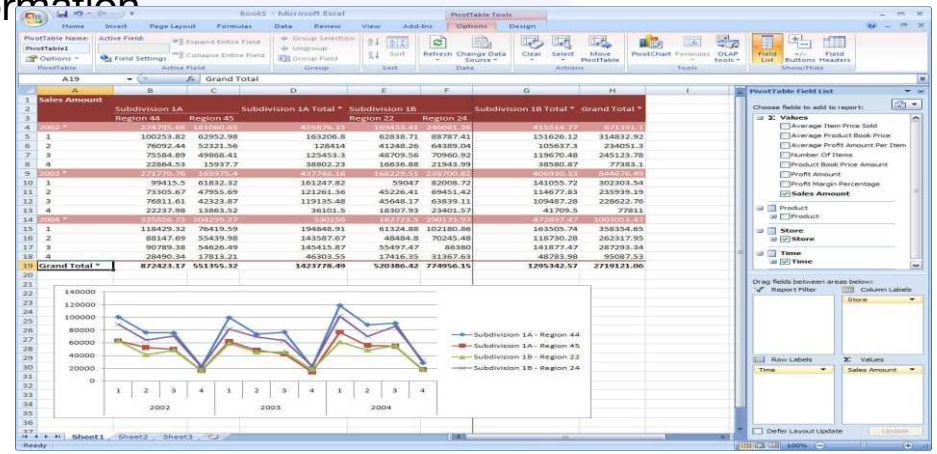

Dynamic Warehouse
The Integrated Stack for System z

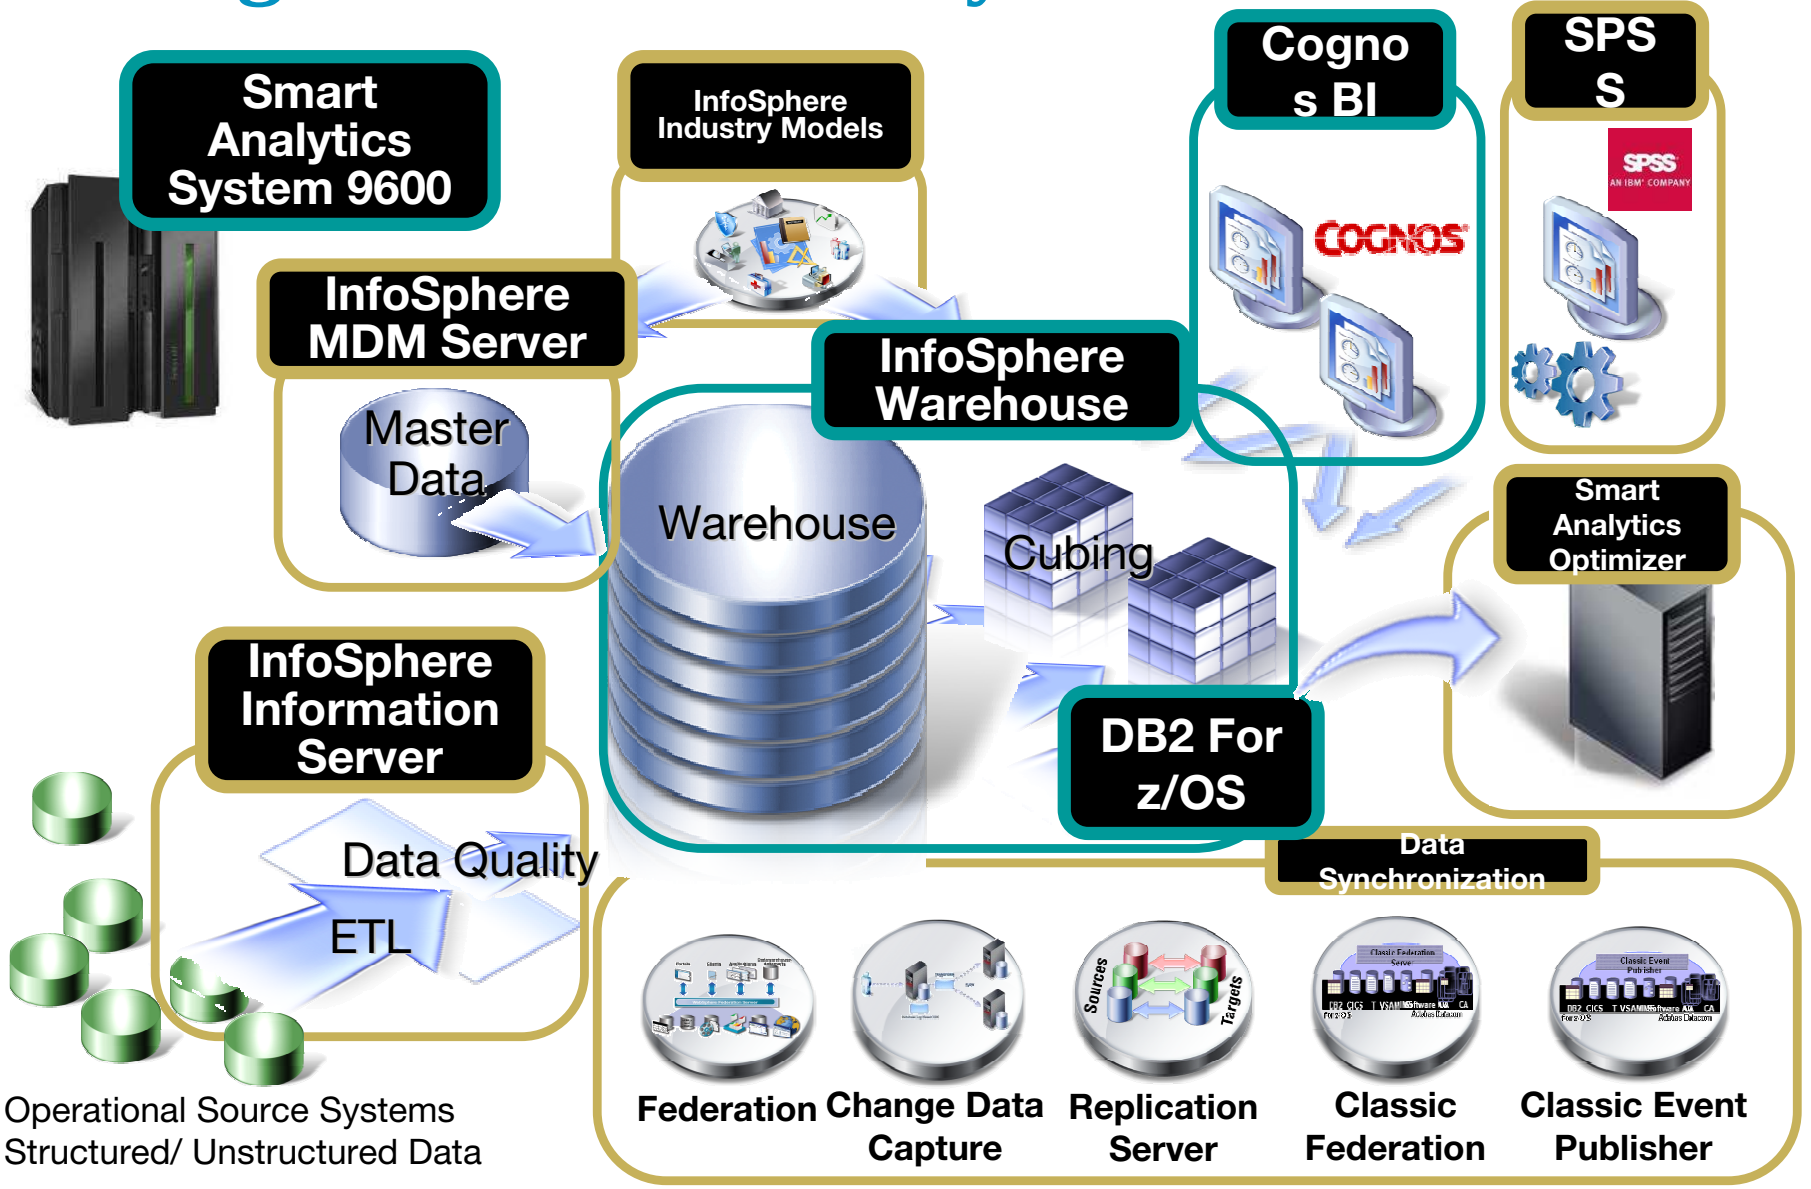

Benefits of a DB2 for z/OS Warehouse

- Blurring of distinction between warehousing applications and OLTP
  - Users demanding reliable access to reliable data from multiple sources
  - Global competitiveness requires 24 x 7
  - Evolving demand for real time or near real time
- Operational data and the ODS together means
  - Reduced complexity
  - Reduced cost
  - Shared processes, tools, procedures
  - Streamlined compliance and security
- Significant DW capabilities in V10 make DB2 competitive
  - DB2 VUE
- zIIP specialty engine allows for IT optimization
- Unprecedent query performance with SAO appliance
- Bundled offering with ISAS 9600
- Better leverage System z skills and investment
- No other solution can match the qualities of service of z